ComDirect

Network Optimization for the Hybrid Workforce

Strategies for Designing and Optimizing Enterprise Networks That Support Both On-Premise and Remote Operations

Published Q3 2024

ComDirect USA

comdirectusa.com

This white paper is intended for informational purposes only. The strategies and frameworks described are general industry guidance and should be adapted to your organization's specific requirements and network environment.

Table of Contents

Chapter 1

The New Network Reality

The enterprise network was designed for a world that no longer exists. Traditional hub-and-spoke architectures assumed that employees, applications, and data all lived inside the same building. Traffic flowed through a central data center, firewalls inspected everything at the perimeter, and bandwidth planning was straightforward.

That model collapsed under the weight of three simultaneous shifts: the mass adoption of cloud applications, the permanent embrace of remote and hybrid work, and the explosion of real-time collaboration tools that demand low latency and high throughput. The network that worked in 2019 is now a bottleneck.

Organizations running hybrid workforces face a unique challenge. They must deliver consistent, high-performance connectivity to employees in corporate offices, home offices, co-working spaces, and on the road, all while maintaining security, visibility, and cost control.

The Hybrid Reality

Industry surveys consistently show that 60-70% of knowledge workers now split their time between office and remote locations. Network architectures must treat remote connectivity as a first-class requirement, not an afterthought bolted onto a campus-centric design.

This white paper provides a practical framework for optimizing your enterprise network to support hybrid operations. It covers SD-WAN deployment, bandwidth management, QoS policies, secure remote access, cloud networking, performance monitoring, and cost optimization.

Chapter 2

Assessing Your Current Network Architecture

Before optimizing, you need an honest assessment of where you stand. Many organizations are running network architectures that evolved organically over years, with layers of legacy equipment, inconsistent configurations, and undocumented dependencies.

Network Discovery and Inventory

Start with a complete inventory of your network infrastructure:

- All routers, switches, firewalls, load balancers, and wireless access points, including model, firmware version, and end-of-support dates
- WAN circuits: carrier, bandwidth, contract terms, and actual utilization
- Internet breakout points: where does traffic exit to the internet, and how much of it backhaults to a central site first?
- VPN concentrators and remote access infrastructure: capacity, concurrent session limits, and geographic distribution

Traffic Flow Analysis

Understanding how traffic actually moves through your network is essential for optimization:

- What percentage of traffic is destined for cloud/SaaS applications vs. on-premise resources?
- Which applications consume the most bandwidth? Are they business-critical or best-effort?
- Where are the latency-sensitive applications (voice, video, real-time collaboration)?
- How much traffic is backhauled through a central data center that could be broken out locally?

Common Finding

Most organizations discover that 60-80% of their WAN traffic is now destined for the internet or cloud services, yet their architecture still forces all of it through a central data center. This single finding often justifies the entire optimization project.

Chapter 3

SD-WAN: The Foundation of Hybrid Connectivity

Software-Defined Wide Area Networking (SD-WAN) has become the foundational technology for hybrid workforce networking. It replaces rigid, expensive MPLS-centric architectures with intelligent, application-aware routing across multiple transport types.

What SD-WAN Actually Does

At its core, SD-WAN decouples the network control plane from the transport layer. This means:

- Traffic can be dynamically routed across MPLS, broadband internet, LTE/5G, or any combination based on real-time performance metrics
- Application-aware policies steer business-critical traffic over the best-performing path, not just the shortest path
- Centralized management provides a single pane of glass for policy configuration, monitoring, and troubleshooting across all sites
- Built-in encryption secures traffic across public internet transports, reducing or eliminating the need for dedicated MPLS circuits

SD-WAN Deployment Models

Choose the deployment model that matches your operational maturity:

- DIY: You purchase SD-WAN appliances and manage them with your own team. Maximum control, maximum operational burden.
- Co-managed: A managed service provider handles day-to-day operations while you retain policy control and visibility. Balances control with operational efficiency.
- Fully managed: The provider owns the entire SD-WAN lifecycle from design through operations. Best for organizations without deep networking staff.

MPLS Reduction Strategy

SD-WAN does not always mean eliminating MPLS entirely. A common pattern is to retain a reduced MPLS backbone for the most latency-sensitive traffic (voice, database replication) while offloading cloud and internet traffic to broadband circuits at each site. This hybrid transport approach typically reduces WAN costs by 30-50% while improving application performance.

Chapter 4

Bandwidth Management and Traffic Shaping

Bandwidth is a finite resource, and hybrid work has made demand less predictable. Offices that were designed for 500 users may now host 200 on any given day, while the other 300 are consuming bandwidth from home connections of varying quality.

Application-Based Bandwidth Allocation

Stop thinking about bandwidth in terms of circuits and start thinking about it in terms of applications:

- Tier 1 (Guaranteed): Voice, video conferencing, and real-time collaboration tools get reserved bandwidth that cannot be consumed by other traffic
- Tier 2 (Priority): Business-critical SaaS applications (CRM, ERP, email) get priority access to remaining bandwidth
- Tier 3 (Best-Effort): General web browsing, software updates, and non-critical cloud storage sync use whatever bandwidth is available
- Tier 4 (Restricted): Personal streaming, social media, and large file downloads are rate-limited or scheduled for off-peak hours

Dynamic Bandwidth Scaling

Static bandwidth allocations waste money. Modern SD-WAN and cloud networking solutions support dynamic scaling:

- Burst capabilities on broadband circuits for temporary demand spikes
- LTE/5G failover that activates automatically when primary circuits are saturated
- Cloud-based internet breakout that scales with the number of users at a given site
- Time-of-day policies that reallocate bandwidth based on usage patterns (e.g., allowing software updates to consume full bandwidth overnight)

Bandwidth Waste

Organizations that implement application-based traffic shaping typically discover that 15-25% of their bandwidth was being consumed by non-business traffic (OS updates, cloud backup sync, streaming media) during peak hours. Rescheduling or rate-limiting this traffic often eliminates the need for bandwidth upgrades.

Chapter 5

Quality of Service (QoS) for Real-Time Applications

Bandwidth is only half the equation. Real-time applications like voice and video are far more sensitive to latency, jitter, and packet loss than they are to raw throughput. A 100 Mbps connection with 200ms of latency and 2% packet loss will deliver a worse video call than a 10 Mbps connection with 20ms latency and zero packet loss.

QoS Fundamentals

QoS policies classify, mark, and prioritize traffic to ensure that latency-sensitive applications get preferential treatment through every hop in the network:

- Classification: Identify traffic by application (deep packet inspection), port, DSCP marking, or source/destination
- Marking: Tag packets with a priority level (DSCP values) so that every device in the path knows how to handle them
- Queuing: Place high-priority traffic in low-latency queues and best-effort traffic in standard queues
- Policing and shaping: Enforce bandwidth limits on lower-priority traffic to protect headroom for real-time applications

Voice and Video Optimization

For unified communications platforms (Teams, Zoom, Webex, etc.):

- Allocate a minimum of 100 kbps per concurrent voice call and 2-4 Mbps per video stream
- Target latency below 150ms one-way, jitter below 30ms, and packet loss below 1%
- Use local internet breakout for cloud-based UC platforms rather than backhauling through a central site
- Deploy session border controllers (SBCs) at the network edge for SIP-based voice traffic

QoS Across the WAN

QoS is only effective end-to-end. Marking packets on the LAN is pointless if the WAN carrier does not honor those markings. Verify that your MPLS provider supports your DSCP scheme. For internet-based transports, SD-WAN provides application-aware path selection that achieves similar results without relying on carrier QoS.

Chapter 6

Secure Remote Access Without VPN Bottlenecks

Traditional VPN architectures were designed for occasional remote access by a small percentage of the workforce. When the entire company went remote, VPN concentrators became chokepoints: oversubscribed, latency-inducing, and operationally fragile.

The Problem with Full-Tunnel VPN

Full-tunnel VPN routes all traffic from the remote user through the corporate data center, even traffic destined for cloud services. This creates:

- Unnecessary latency as cloud-bound traffic hairpins through the data center
- Bandwidth congestion at the VPN concentrator and data center internet egress
- A single point of failure that can disconnect the entire remote workforce
- Poor user experience that drives employees to use personal devices and shadow IT

Split Tunneling Done Right

Split tunneling routes only corporate-bound traffic through the VPN while allowing cloud and internet traffic to break out directly from the user's local connection. Modern split tunneling is application-aware:

- Traffic to trusted SaaS platforms (Microsoft 365, Salesforce, etc.) goes direct
- Traffic to internal applications and sensitive data stores routes through the VPN
- DNS-based policies prevent data exfiltration through unauthorized cloud services

Zero Trust Network Access (ZTNA)

ZTNA replaces VPN entirely for many use cases. Instead of granting network-level access, ZTNA brokers application-level access based on user identity, device posture, and risk signals. Benefits include:

- No inbound firewall rules or exposed VPN concentrators
- Per-application access instead of broad network access
- Consistent experience whether the user is on-premise, at home, or on the road
- Reduced attack surface compared to traditional VPN

Migration Path

Most organizations adopt ZTNA incrementally: start with web-based applications (easiest to broker), then extend to thick-client applications, and finally retire the legacy VPN once all critical applications are covered. This typically takes 6-12 months.

Chapter 7

Cloud Networking and Direct Connect Strategies

As workloads move to the cloud, the network path between your users and those workloads becomes critical. Relying on the public internet for access to production cloud environments introduces unpredictable latency, packet loss, and security risks.

Direct Connect and ExpressRoute

Major cloud providers offer dedicated, private connectivity options:

- AWS Direct Connect: Dedicated 1-100 Gbps connections bypassing the public internet
- Azure ExpressRoute: Private peering with Microsoft's global network for Azure and Microsoft 365
- Google Cloud Interconnect: Dedicated or partner connections to Google Cloud regions

These services provide consistent latency, higher throughput, and reduced data transfer costs compared to internet-based access. They are essential for latency-sensitive workloads like databases, real-time analytics, and VDI.

Cloud Exchange and Colocation

If you operate from a colocation facility, cloud exchange services (Equinix Fabric, Megaport, etc.) allow you to establish private connections to multiple cloud providers from a single physical location. This simplifies multi-cloud networking and provides redundant paths.

SaaS Optimization

For SaaS applications, direct connect is not always necessary or cost-effective. Instead:

- Use local internet breakout at each site so users reach the nearest SaaS edge node
- Deploy a cloud access security broker (CASB) for visibility and control without backhauling
- Leverage SD-WAN application-aware routing to select the best path to SaaS endpoints
- Monitor SaaS performance with digital experience monitoring (DEM) tools to catch provider-side issues

Chapter 8

Wi-Fi 6/6E and the Modern Office

The hybrid office demands more from wireless infrastructure than ever before. With fewer wired desks and more employees moving between meeting rooms, collaboration spaces, and hot desks, Wi-Fi is the primary network connection for most office workers.

Why Wi-Fi 6/6E Matters

Wi-Fi 6 (802.11ax) and Wi-Fi 6E introduce capabilities specifically designed for high-density, mixed-use environments:

- OFDMA (Orthogonal Frequency Division Multiple Access): Allows a single access point to serve multiple clients simultaneously, reducing latency in crowded spaces
- BSS Coloring: Reduces co-channel interference between adjacent access points
- Target Wake Time (TWT): Improves battery life for IoT devices by scheduling transmission windows
- 6 GHz band (Wi-Fi 6E): Provides 1,200 MHz of clean, uncongested spectrum for high-bandwidth applications

Design Considerations for Hybrid Offices

- Survey and design for peak hybrid occupancy, not average. If 80% of staff comes in on Tuesdays, that is your design target.
- Prioritize meeting rooms and collaboration spaces with dedicated access points and higher density coverage
- Separate IoT devices (printers, displays, sensors) onto a dedicated SSID and VLAN to prevent them from consuming capacity on the primary network
- Deploy wireless intrusion prevention to detect and contain rogue access points

Cabling Still Matters

Wi-Fi 6/6E access points require multi-gigabit uplinks (2.5G or 5G Ethernet) and PoE++ power. Many organizations upgrade their access points only to discover that their existing Cat 5e cabling and legacy switches cannot deliver the required power and throughput. Budget for switching and cabling upgrades alongside AP deployment.

Chapter 9

Performance Monitoring for Distributed Teams

You cannot optimize what you cannot see. In a hybrid environment, network performance monitoring must extend beyond the corporate LAN to encompass home networks, ISP last-mile connections, cloud paths, and SaaS application performance.

Digital Experience Monitoring (DEM)

DEM tools measure the end-to-end experience from the user's device to the application and back. They provide visibility into:

- Device-level metrics: CPU, memory, Wi-Fi signal strength, and DNS resolution time on the endpoint
- Network path metrics: Latency, jitter, packet loss, and hop-by-hop traceroute for every connection
- Application-level metrics: Page load time, API response time, error rates, and transaction completion
- Synthetic monitoring: Simulated user transactions that detect problems before real users are affected

Network Performance Baselines

Establish baselines for key metrics so you can detect degradation before users complain:

- LAN latency: < 1ms between client and distribution switch
- WAN latency: < 50ms between branch sites (varies by geography)
- Internet latency: < 20ms to major cloud provider edge nodes from local breakout
- Packet loss: < 0.1% on any segment under normal load
- Wi-Fi client experience: > -65 dBm signal strength, < 10ms airtime latency

Alerting and Escalation

Configure tiered alerts based on severity and business impact:

- Warning: Metric exceeds baseline by 20%, triggers monitoring team review
- Critical: Metric exceeds baseline by 50% or crosses a hard threshold (e.g., > 200ms latency), triggers immediate investigation
- Outage: Complete loss of connectivity or application availability, triggers incident response

Chapter 10

Redundancy, Failover, and Business Continuity

A network optimized for performance but lacking redundancy is a single point of failure. Hybrid workforces depend on the network for everything: communication, collaboration, access to applications, and access to data. Downtime is not an inconvenience; it is a full work stoppage.

WAN Redundancy

- Dual-carrier internet at every site, using different physical paths where possible (fiber + cable, or fiber + LTE)
- SD-WAN automatic failover that switches traffic to the backup circuit within seconds, transparent to users
- LTE/5G as a tertiary failover for critical sites, with pre-provisioned SIM cards and tested connectivity

DNS and Cloud Service Redundancy

- Use multiple DNS providers (e.g., primary cloud DNS with a secondary provider) to prevent DNS-related outages
- Configure health-check-based DNS failover for public-facing services
- Distribute cloud workloads across availability zones and, for critical services, across regions

Remote Worker Resilience

Home network failures are outside your control, but you can mitigate their impact:

- Provide company-issued LTE/5G hotspots for critical roles
- Ensure mobile applications support offline modes for essential workflows
- Pre-configure failover to cellular on company-issued laptops with embedded LTE
- Document co-working space and backup location options for extended home outages

Testing Is Non-Negotiable

Redundancy that has never been tested is not redundancy. Schedule quarterly failover tests for every critical circuit and service. Simulate primary link failure during business hours and measure failover time, user impact, and restoration. Document results and fix any gaps before they become real incidents.

Chapter 11

Cost Optimization and Vendor Management

Network costs are often the second-largest line item in the IT budget after headcount. Hybrid workforce optimization creates opportunities to reduce costs while improving performance, but only if you approach vendor management strategically.

Circuit Cost Reduction

- Replace expensive MPLS circuits with broadband + SD-WAN at branch sites. Retain MPLS only where dedicated SLAs are required.
- Negotiate multi-year broadband contracts with bandwidth tiers that match actual usage, not peak theoretical demand
- Consolidate carriers where possible to increase purchasing leverage, but maintain at least two carriers per site for redundancy
- Audit circuits monthly for utilization. Downgrade or terminate circuits running below 30% sustained utilization.

Cloud Networking Costs

- Direct connect services charge per port-hour and per GB of data transfer. Right-size your port speed and monitor transfer volumes.
- Use cloud provider cost calculators to compare direct connect vs. VPN vs. internet egress costs for your traffic volumes
- Implement data transfer policies that minimize cross-region and cross-cloud data movement

Vendor Consolidation

Many organizations run separate vendors for SD-WAN, firewall, wireless, and monitoring. Converged platforms that combine multiple functions can reduce licensing costs, simplify operations, and improve integration. Evaluate whether a SASE (Secure Access Service Edge) platform could consolidate your SD-WAN, ZTNA, CASB, and firewall-as-a-service into a single vendor relationship.

Total Cost of Ownership

When comparing solutions, include operational costs (staff time for management, troubleshooting, and vendor coordination), not just licensing and circuit fees. A cheaper product that requires twice the operational effort is not actually cheaper.

Chapter 12

Building Your Network Optimization Roadmap

Network optimization is an ongoing discipline, not a one-time project. Build a phased roadmap that delivers quick wins early while setting the foundation for long-term transformation.

Phase 1: Visibility and Assessment (Months 1-2)

- Complete network inventory and traffic flow analysis
- Deploy digital experience monitoring for baseline metrics
- Identify top bandwidth consumers and latency-sensitive applications
- Audit current circuit contracts, utilization, and costs

Phase 2: Quick Wins (Months 3-5)

- Implement application-based QoS policies for voice and video
- Enable local internet breakout for cloud and SaaS traffic at branch sites
- Deploy split tunneling or ZTNA for remote workers
- Terminate or downgrade underutilized circuits

Phase 3: Architecture Transformation (Months 6-10)

- Deploy SD-WAN across branch sites with dual-transport redundancy
- Implement direct connect for primary cloud platforms
- Upgrade wireless infrastructure to Wi-Fi 6/6E in high-density locations
- Establish micro-segmentation for IoT and guest networks

Phase 4: Continuous Optimization (Ongoing)

- Review performance baselines and adjust thresholds quarterly
- Conduct failover testing on a regular schedule
- Renegotiate vendor contracts as traffic patterns evolve
- Evaluate emerging technologies (Wi-Fi 7, private 5G, SASE convergence)

About ComDirect USA

ComDirect USA is a full-service technology partner specializing in enterprise infrastructure, cybersecurity, cloud architecture, and managed IT services. With offices in Irvine, California, and Meridian, Idaho, we serve organizations across healthcare, legal, financial services, manufacturing, and logistics.

Contact Us

Phone: 949.208.8800 | Toll Free: 800.733.7030

Email: info@comdirectusa.com

Web: comdirectusa.com